arXiv:1402.5090v3 [stat.ME] 31 Mar 2014

MAD Bayes for Tumor Heterogeneity – Feature Allocation with Exponential Family Sampling

Yanxun Xu

Division of Statistics and Scientific Computing, The University of Texas at Austin, Austin, TX

Peter Müller *

Department of Mathematics, The University of Texas at Austin, Austin, TX

Yuan Yuan

SCBMB, Baylor College of Medicine, Houston, TX

Kamalakar Gulukota

Center for Biomedical Informatics, NorthShore University HealthSystem, Evanston, IL

Yuan Ji *

Center for Biomedical Informatics, NorthShore University HealthSystem, Evanston, IL

Department of Health Studies, The University of Chicago, Chicago, IL

*Address for Correspondence: Department of Mathematics UT Austin 1, University Station, C1200, Austin, TX 78712 USA. E-mail: pmueller@math.utexas.edu. jiyuan@uchicago.edu.

Abstract

We propose small-variance asymptotic approximations for the inference of tumor heterogeneity (TH) using next-generation sequencing data. Understanding TH is an important and open research problem in biology. The lack of appropriate statistical inference is a critical gap in existing methods that the proposed approach aims to fill. We build on a hierarchical model with an exponential family likelihood and a feature allocation prior. The proposed approach generalizes similar small-variance approximations proposed by Kulis and Jordan (2012) and Broderick et.al (2012) for inference with Dirichlet process mixture and Indian buffet prior models under normal sampling. We show that the new algorithm can successfully recover latent structures of different subclones and is also magnitude faster than available Markov chain Monte Carlo samplers, the latter often practically infeasible for high-dimensional genomics data. The proposed approach is scalable, simple to implement and benefits from the flexibility of Bayesian nonparametric models. More importantly, it provides a useful tool for the biological community for estimating cell subtypes in tumor samples. R code is available on http://www.ma.utexas.edu/users/yxu/.

KEY WORDS: Bayesian nonparametric; Bregman divergence; Feature allocation; Indian buffet process; Next-generation sequencing; Tumor heterogeneity.

1 Introduction

1.1 MAD-Bayes

We propose a generalization of the MAD (MAP-based asymptotic derivations) Bayes approach of Broderick et al. (2012b) to latent feature models beyond the conjugate normal-normal setup. The model is developed for inference on tumor heterogeneity (TH), when the sampling model is a binomial distribution for observed short reads counts for single nucleotide variants (SNVs) in next-generation sequencing (NGS) experiments.

The proposed model includes a Bayesian non-parametric (BNP) prior. BNP models are characterized by parameters that live on an infinite-dimensional parameter space, such as unknown mean functions or unknown probability measures. Related methods are widely used in a variety of machine learning and biomedical research problems, including clustering, regression and feature allocation. While BNP methods are flexible from a modeling perspective, a major limitation is the computational challenge that arises in posterior inference with large-scale problems and big data. Posterior inference in highly structured models is often implemented by Markov chain Monte Carlo (MCMC) simulation (Liu, 2008, for example) or variational inference such as expectation-maximization (EM) algorithm (Dempster et al., 1977). However, neither approach scales effectively to high-dimensional data. As a result, simple ad-hoc methods, such as K-means (Hartigan and Wong, 1979), are still preferred in many large-scale applications. For example, K-means clustering is often preferred over full posterior inference in model-based clustering, such as Dirichlet process (DP) mixture models (see, for example, Ghoshal (2010), for a review). DP mixture models some of the most widely used BNP models.

Despite the simplicity and scalability, K-means has some known shortcomings. First, the K-means algorithm is a rule-based method. The output is an ad-hoc point estimate of the unknown partition. There is no notion of characterizing uncertainty, and it is difficult to coherently embed it in a larger model. Second, the K-means algorithm requires a fixed number of clusters, which is not available in many applications. An ideal algorithm should combine the scalability of K-means with the flexibility of Bayesian nonparametric models. Such links between non-probabilistic (i.e., rule-based methods like K-means) and probabilistic approaches (e.g., posterior MCMC or the EM algorithm) can sometimes be found by using small-variance asymptotics. For example, the EM algorithm for a mixture of Gaussian model becomes the K-means algorithm as the variance of the Gaussians tend to zero (Hastie et al., 2001). In general, small-variance asymptotics can offer useful alternative approximate

implementations of inference for large-scale Bayesian nonparametric models, exploiting the fact that corresponding non-probabilistic models show advantageous scaling properties.

Using small-variance asymptotics, Kulis and Jordan (2011) showed how a K-means-like algorithm can approximate posterior inference for Dirichlet process (DP) mixtures. Broderick et al. (2012b) generalized the approach by developing small-variance asymptotics to MAP (maximum a posteriori) estimation in feature allocation models with Indian buffet process (IBP) priors (Griffiths and Ghahramani, 2006; Teh et al., 2007). Similar to the K-means algorithm, they proposed the BP (beta process)-means algorithm for feature learning. Both approaches are restricted to normal sampling and conjugate normal priors, which facilitates the asymptotic argument and greatly simplifies the computation. However, it is not immediately generalizable to other distributions, preventing their methodology from being applied to non-Gaussian data. The application that motivates the current paper is a typical example. We require posterior inference for a feature allocation model with a binomial sampling model.

1.2 Tumor heterogeneity

The proposed methods are motivated by an application to inference for tumor heterogeneity (TH). This is a highly important and open research problem that is studied by many cancer researchers (Gerlinger et al., 2012; Landau et al., 2013). In the literature over the past five years a consensus has emerged that tumor cells are heterogenous, both within the same biological tissue sample and between different samples. A tumor sample typically comprises an admixture of subtypes of different cells, each possessing a unique genome. We will use the term "subclones" to refer to cell subtypes in a biological sample. Genotypic differences (differences in DNA base pairs) between subclones and proportions of each subclone in a sample will provide ground-breaking new information that is useful for cancer diagnosis and prognosis. However, inference and statistical modeling are challenging and few solutions

exist.

Genotypic differences between subclones do not occur frequently. They are often restricted to single nucleotide variations (SNVs). For the purpose of this study, we consider subclones in multiple biological samples. We assume that cells in each sample are mosaic, consisting of multiple subclones. Each subclone is characterized by a set of somatically acquired SNVs. More importantly, we regard two subclonal genomes the same if they possess identical genotypes on the SNVs, regardless of the rest of genome. In other words, we do not insist that the whole genomes of any two cells must be identical in order to call them subclonal. This is because 1) two cells from different tumor samples almost never share the same whole genome (due to the genomic mosaicism in the species) and 2) only part of the genome is assumed to be related to the disease. Hereafter, we refer to a subclone as cells sharing the same genotypes on selected SNVs. We say that cells in a sample are mosaic, with co-existing subclones possessing different genotypes on the SNVs.

To study TH, we use NGS data that record the numbers of read counts that overlap with SNVs for each of a number of tumor samples. For each SNV, there are many overlapping short reads, some of which possess a variant sequence while others include the reference sequence. In Figure 1, five reads are mapped to the SNV with reference sequence "A". Two short reads are variants bearing a sequence "C". We define variant allele fraction (VAF) as the proportion of short reads bearing a variant genotype among all the short reads mapped to a SNV. In Figure 1, the VAF is 2/5 for that SNV. In Section 2, we will model the observed variant read counts by a mixture of latent subclones, which in turn are defined by a pattern of present and absent SNVs. Assuming that each sample is composed of some proportions of these subclones, we can then fit the observed VAFs across SNVs in each sample. Formally, modeling involves binomial sampling models for the observed counts with mixture priors for the binomial success probabilities. The mixture is over an (unknown) number C of (latent) subclones, which in turn are represented as a binary matrix $\boldsymbol{Z}$, with columns, $c = 1, \ldots, C$,

corresponding to subclones and rows, $s = 1, \ldots, S$, corresponding to SNVs. The entries $z_{sc} \in \{0, 1\}$ are indicators for variant alleles s appearing in subclone c. That is, each column of indicators defines a subclone. Column c spells out the genotypes (variant or not) of the SNVs in the c-th subclone.

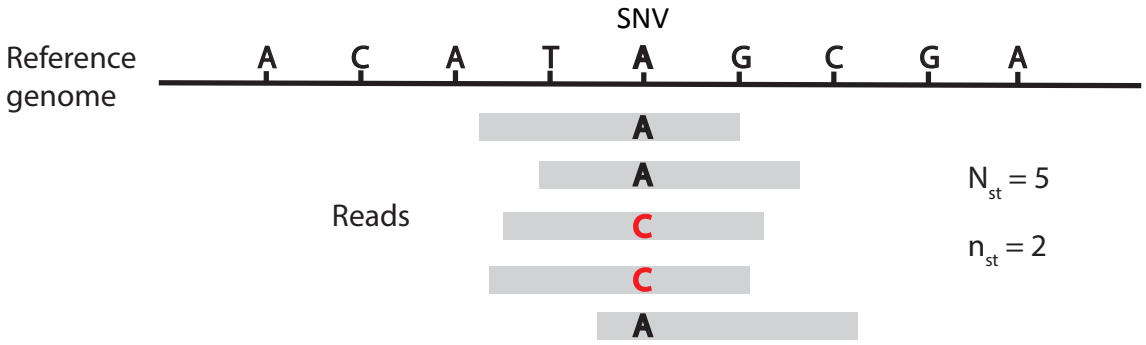

Figure 1: An illustration of SNV and read mapping. The figure shows that five short reads are mapped to the SNV location. The reference genotype is A. Among the five reads, two reads have a variant genotype C. So $N_{st} = 5$, $n_{st} = 2$ and the observed VAF is 2/5.

1.3 Main contributions

A key element of the proposed model is the prior on the binary matrix $\boldsymbol{Z}$. We recognize the problem as a special case of a feature allocation problem and use an Indian buffet process prior for $\boldsymbol{Z}$. In the language of feature allocation models (and the traditional metaphor that is used to describe the Indian buffet process prior), the subclones are the features (or dishes) and the SNVs are experimental units (or customers) that select features. Each tumor sample consists of an unknown proportion of these subclones. Lee et al. (2013) used a finite feature allocation model to describe the latent structure of possible subclones. The model is restricted to a fixed number of subclones. In practice, the number of subclones is unknown. A possible way to generalize to an unknown number of latent features is to define

a transdimensional MCMC scheme, such as a reversible jump (RJ) algorithm (Green, 1995). However, it is difficult to implement a practicable RJ algorithm.

An attractive alternative is to generalize the BP-means algorithm of Broderick et al. (2012b) beyond the Gaussian case. Inspired by the connection between Bregman divergences (Bregman, 1967) and exponential families, we propose a MAP-based small-variance asymptotic approximation for any exponential family likelihood with the IBP as a feature allocation prior. We call the proposed approach the FL-means algorithm, where FL stands for "feature learning". The FL-means algorithm is scalable, simple to implement and benefits from the flexibility of Bayesian nonparametric models. It improves the speed of computation by thousand-folds for big data.

This paper proceeds as follows. In Section 2, we introduce the Bayesian feature allocation model to describe tumor heterogeneity. Section 3 elaborates the FL-means algorithm and proves its convergence property. We examine the performance of FL-means through simulation studies in Section 4. In Section 5, we apply the FL-means algorithm to the tumor heterogeneity examples. Finally, we conclude with a brief discussion in Section 6.

2 A Model for Tumor Heterogeneity

We consider data sets from NGS experiments. The data sets record the observed read counts for S SNVs in T tumor samples. Let $\boldsymbol{n}$ and $\boldsymbol{N}$ denote $S \times T$ matrices with N_{st} denoting the total number of reads overlapping with SNV s in tissue sample t, and n_{st} denoting the number of variant reads among those N_{st} reads. The ratio n_{st}/N_{st} is then the observed VAF. Figure 1 provides an illustration. We assume a binomial sampling model for the counts n_{st}, given by

$$p(n_{st} \mid N_{st}, p_{st}) = \text{Bin}(N_{st}, p_{st}),$$

where p_{st} is the expected VAF, $p_{st} = E(n_{st}/N_{st})$. Conditional on N_{st} and p_{st}, the observed counts n_{st} are independent across s and t. The likelihood becomes

$$p(\boldsymbol{n} \mid \boldsymbol{N}, \boldsymbol{p}) = \prod_{s=1}^{S} \prod_{t=1}^{T} \binom{N_{st}}{n_{st}} p_{st}^{n_{st}} (1 - p_{st})^{(N_{st} - n_{st})}, \tag{1}$$

where $\boldsymbol{n} = [n_{st}]$, $\boldsymbol{N} = [N_{st}]$ and $\boldsymbol{p} = [p_{st}]$ are $(S \times T)$ matrices. The binomial success probabilities p_{st} are modeled in terms of latent subclones which we introduce next. Recall that $\{z_{sc} = 1\}$ indicates that the genotype at SNV s is a variant (different from the reference genotype) for subclone c. In addition, we introduce a background subclone $c = 0$ for which all $\{z_{s0} = 1\}$, i.e., all the SNVs are variants for this subclone. The background subclone captures experimental and data processing noise, such as the fraction of reads that are mapped with low quality or error. Figure 2 illustrates the binary latent matrix $\boldsymbol{Z} = \{z_{sc}\}$ with $C = 4$ subclones (columns) and $S = 8$ SNVs (rows). A shaded cell indicates $z_{sc} = 1$. For instance, SNV 2 occurs in two subclones $c = 1$ and 3, SNV 3 occurs only in subclone $c = 3$.

Assuming that each sample is composed of proportions w_{tc} of the subclones, $c = 0, \ldots, C$, we then represent p_{st} as

$$p_{st} = w_{t0} p_0 + \sum_{c=1}^{C} w_{tc} z_{sc}, \tag{2}$$

where p_0 is the relative frequency of a SNV in the background and $w_{tc} \in (0, 1), \sum_{c=0}^{C} w_{tc} = 1$. Note that the use of p_0 implies that $z_{s0} = 1$ is interpreted as partial presence of SNV s.

Equation (2) is a key model assumption. It allows us to deduce the unknown subclones from a decomposition of the expected VAF p_{st} as a weighted sum of genotype calls z_{sc} with weights being the proportions of subclones. In other words, the expected VAF is contributed by those subclones with variant genotypes, weighted by the subclonal prevalences. The

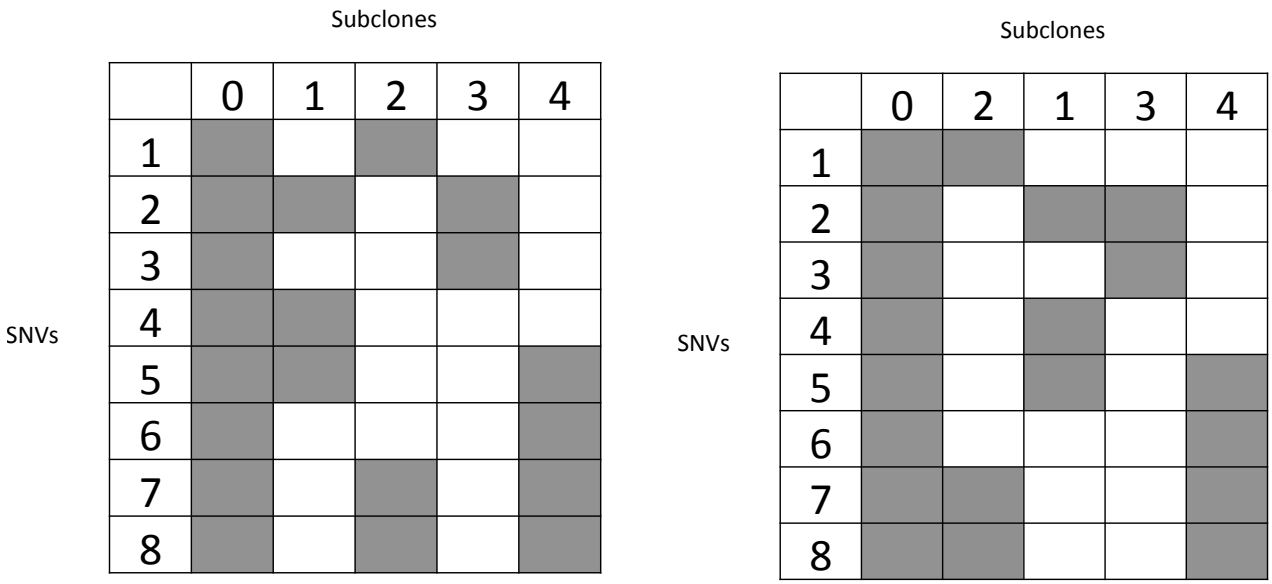

(a) $\boldsymbol{Z}$ (b) left-ordered form of $\boldsymbol{Z}$

Figure 2: The binary matrix $\boldsymbol{Z}$ and the left-ordered form of $\boldsymbol{Z}$, with subclones (latent features) $c = 1, \ldots, C = 4$, in the columns and $S = 8$ SNVs in the rows. The shaded cell represents that a SNV feature occurs in a subclone, that is $z_{sc} = 1$. In addition, we show the background subclone $c = 0$ with all $z_{s0} = 1$ in the first column. Matrix in (b) is the left-ordered form of the matrix in (a).

subclones without variant genotype on SNV s do not contribute to VAF for s since all the short reads generated from those subclones are normal reads. For example, in Figure 2 the variant reads mapped to SNV 2 should come from subclones 1 and 3, but not from 2 or 4. Another critical assumption is the prior model $p(\boldsymbol{Z})$. We use the Indian buffet process (IBP) prior. See, for example Griffiths and Ghahramani (2006); Thibaux and Jordan (2007) or Teh et al. (2007) for a discussion of the IBP. It is customary to restrict $\boldsymbol{Z}$ to left-ordered form. Let $f_c = \min\{s : \ z_{sc} = 1\}$ denote the first SNV that is included in a given subclone. Under the left ordered form, columns are ordered such that $1 = f_1 \leq f_2 \ldots \leq f_C$. This ordering is natural under the metaphor that gives rise to the name of the IBP. If z_{sc} are indicators for customer s selecting dish c, then the left-ordered form simply implies that dishes are sorted by order of selection as customers $s = 1, \ldots, S$ enter the restaurant. Let $m_{sc} = \sum_{s'<s} z_{s'c}$ denote the column sums up to row $s-1$, that is, the number of SNVs, $s' < s$, that define subclone c,

and let m_c denote the total column sum of subclone c. Also let $C_s^+ = \sum_{c=1}^{C} I(f_c = s)$ denote the number of subclones for which s is the first SNV to define the subclone. In the original IBP metaphor, that is the number of new dishes (subclones) that are selected by customer (SNV) s. Let $C_s = \sum_{s'=1}^{s} C_{s'}^+$ denote the total number of subclones selected by the first s SNVs, and let $C = C_S$ be the total number of latent subclones. Note that the cumulative sum C_s includes the s-th SNV, while m_{sc} does not. Finally, let $H_S = \sum_{s=1}^{S} 1/s$. Then the IBP prior for a left ordered $(S \times C)$ binary matrix can be defined by the following generative model. For each SNV, $s = 1, \ldots, S$, introduce $C_s^+ \sim \text{Poi}(\gamma/s)$ new subclones, and assign $z_{sc} = 1$ and $z_{s',c} = 0, \forall s' < s$, for the new subclones $c = C_{s-1}+1, \ldots, C_{s-1}+C_s^+ = C_s$. For the earlier subclones, $c = 1, \ldots, C_{s-1}$, mutation s is included with probability $p(z_{sc} = 1) = \frac{m_{sc}}{s}$ (note the s, rather than $s-1$ in the denominator).

After removing the order constraint, that is, with uniform permutation of the columns, we get the IBP prior for a binary matrix, *without left order constraint*,

$$p_{IBP}(\boldsymbol{Z} \mid \gamma) = \frac{\gamma^C \, e^{-\gamma H_S}}{C!} \prod_{c=1}^{C} \frac{(S - m_c)!(m_c - 1)!}{S!}. \tag{3}$$

See, Griffiths and Ghahramani (2006) or Broderick et al. (2012a) for a derivation of (3). In the application to TH, index the columns of the latent matrix $\boldsymbol{Z}$ by $c = 0, 1, \ldots, C$, with $c = 0$ indicating the first column corresponding to the hypothetical background subclone. We assume (3) as prior for columns $c = 1$ through C of the binary matrix $\boldsymbol{Z}$. Let $\boldsymbol{Z}_{-0}$ denote the binary matrix with column $c = 0$ removed. We define

$$p(\boldsymbol{Z}) = p_{IBP}(\boldsymbol{Z}_{-0}). \tag{4}$$

The first column of $\boldsymbol{Z}$ is fixed with all 1's, corresponding to the hypothetical background type.

Finally, the model is completed with a prior on $\boldsymbol{w}_s = (w_{s0}, \ldots, w_{sC})$, $s = 1, \ldots, S$. We assume independent Dirichlet priors, independent across SNV's s, $\boldsymbol{w}_s \sim \text{Dir}(a_0, \ldots, a_C)$, using a common value $a_1 = \ldots = a_C = a$ and distinct a_0.

In summary, the hierarchical model factors as

$$p(\boldsymbol{Z}, \boldsymbol{w}, \boldsymbol{n} \mid \boldsymbol{N}) = \underbrace{p(\boldsymbol{Z})}_{(4)} p(\boldsymbol{w} \mid \boldsymbol{Z}) \underbrace{p(\boldsymbol{n} \mid \boldsymbol{p}, \boldsymbol{N})}_{(1)}. \tag{5}$$

Recall that p_{st} is specified in (2) as a deterministic function of $\boldsymbol{Z}$ and $\boldsymbol{w}$, that is $p_{st} = p_{st}(\boldsymbol{w}, \boldsymbol{Z})$. The joint posterior $p(\boldsymbol{Z}, \boldsymbol{w} \mid \boldsymbol{N}, \boldsymbol{n})$, and thus the desired inference on TH are well defined by (5). However, practically useful inference requires summaries, such as the MAP (maximum a posteriori) estimate or efficient posterior simulation from (5), which could be used to compute (most) posterior summaries. Unfortunately, both, MAP estimation and posterior simulation are difficult to implement here. We therefore extend the approach proposed in Broderick et al. (2012b), who define small-variance asymptotics for inference under an IBP prior using a standard mixture of Gaussian sampling model. We develop a similar approach for the binomial model (1), and it applies to any other exponential family sampling model.

3 A MAD Bayes Algorithm for TH

3.1 Bregman divergence and the scaled binomial distribution

We define small-variance asymptotics for any exponential family sampling model, including in particular the binomial sampling model (1). The idea is to first rewrite the exponential family model in the canonical form as a function of a generalized distance between the random variable and the mean vector. We use Bregman divergence to do this. In the canonical form it is then possible to define a natural scale parameter which becomes the target of the desired

asymptotic limit. Finally, we will recognize the log posterior as approximately equal to a K-means type criterion. The latter will allow fast and efficient evaluation of the MAP. Repeat computations with different starting values finds a set of local modes, which are used to summarize the posterior distribution. We start with a definition of Bregman divergence.

Definition Let $\phi : \mathcal{S} \to \mathbb{R}$ be a differentiable, strictly convex function defined on a convex set $\mathcal{S} \subseteq \mathbb{R}^n$. The Bregman divergence (Bregman, 1967) for any points $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ is defined as

$$d_\phi(\boldsymbol{x}, \boldsymbol{y}) = \phi(\boldsymbol{x}) - \phi(\boldsymbol{y}) - \langle \boldsymbol{x} - \boldsymbol{y}, \nabla\phi(\boldsymbol{y}) \rangle,$$

where $\langle \cdot, \cdot \rangle$ represents the inner product and $\nabla\phi(\boldsymbol{y})$ is the gradient vector of ϕ.

In words, d_ϕ is defined as the increment $\{\phi(\boldsymbol{x}) - \phi(\boldsymbol{y})\}$ beyond a linear approximation with the tangent in $\boldsymbol{y}$. The Bregman divergence leads to a large number of useful divergences as special cases, such as squared loss distance, KL-divergence, logistic loss, etc. For instance, if $\phi(\boldsymbol{x}) = \langle \boldsymbol{x}, \boldsymbol{x} \rangle$, $d_\phi(\boldsymbol{x}, \boldsymbol{y}) = ||\boldsymbol{x} - \boldsymbol{y}||^2$ is the squared Euclidean distance.

Banerjee et al. (2005) show that there exists a unique Bregman divergence corresponding to every regular exponential family including binomial distribution. Specifically, defining the natural parameter $\eta = \log(\frac{p}{1-p})$, we rewrite the probability mass function of $n \sim \text{Bin}(N, p)$ in the canonical form, given by

$$p(n \mid \eta, \psi) = \exp\{n\eta - \psi(\eta) - h_1(n)\}, \tag{6}$$

where $\psi(\eta) = N \log(1 + e^\eta)$ and $h_1(n) = -\log \binom{N}{n}$. Under (6) the mean and variance of n can be written as a function of η, given by

$$\mu = \nabla\psi(\eta) = N\frac{e^\eta}{1 + e^\eta} \quad \text{and} \quad \sigma^2 = \nabla^2\psi(\eta) = N\frac{e^\eta}{(1 + e^\eta)^2}. \tag{7}$$

We introduce a rescaled version of the likelihood by a power transformation of the kernel of (6), that is by scaling the first two terms in the exponent, replacing η by $\widetilde{\eta} = \beta\,\eta$ and $\psi(\eta)$ by $\beta\psi(\widetilde{\eta}/\beta)$. Let $\widetilde{p}(n \mid \widetilde{\eta}, \widetilde{\psi})$ denote the power transformed model. A quick check of (7) shows that the mean remains unchanged, $\mu = \nabla\psi(\eta) = \nabla\widetilde{\psi}(\widetilde{\eta})$. That is, $\widetilde{p}(\cdot)$ is a rescaled, tightened version of $p(\cdot)$.

The rescaled model can be elegantly interpreted when we rewrite (6) as a function of Bregman divergence for suitably chosen $\phi(n, \mu)$. The rescaled version arises when replacing d_ϕ by $\beta\, d_\phi(n, \mu)$. Let

$$\phi(n) = n\log(\frac{n}{N}) + (N-n)\log(\frac{N-n}{N}).$$

The Bregman divergence is $d_\phi(n, \mu) = \phi(n) - \phi(\mu) - (n-\mu)\nabla\phi(\mu)$, with which the binomial distribution can be expressed as

$$p(n \mid \eta, \psi) = p(n \mid \mu) = \exp\{-d_\phi(n, \mu)\} f_\phi(n), \tag{8}$$

where $f_\phi(n) = \exp\{\phi(n) - h_1(n)\}$. The derivation of (8) is shown in the Supplementary Material A.

Equation (8) facilitates the introduction of a scale parameter for a small-variance asymptotic argument, which we show next. Consider a scaled binomial distribution with $\tilde{\eta} = \beta\eta$ and $\tilde{\psi}(\tilde{\eta}) = \beta\psi(\tilde{\eta}/\beta)$. There exists a unique scaled binomial distribution in the canonical form (6), $\widetilde{p}(n \mid \widetilde{\eta}, \tilde{\psi})$, and the mean and variance of the scaled binomial is given by

$$\begin{aligned} \widetilde{\mu}(\widetilde{\eta}) &= \nabla_{\widetilde{\eta}}\tilde{\psi}(\tilde{\eta}) = \beta\nabla_{\widetilde{\eta}}\psi(\tilde{\eta}/\beta) = \mu(\eta) \\ \widetilde{\sigma}^2(\widetilde{\eta}) &= \nabla^2_{\widetilde{\eta}}\widetilde{\psi}(\widetilde{\eta}) = \beta\nabla_{\widetilde{\eta}}\{\nabla_{\widetilde{\eta}}\psi(\tilde{\eta}/\beta)\} = \sigma^2(\eta)/\beta. \end{aligned}$$

The important feature of this scaled binomial model is that $\widetilde{\sigma}^2 \to 0$ as $\beta \to \infty$, while $\widetilde{\mu}$ remains unchanged. Denoting $\widetilde{\phi} = \beta\phi$, we show the Bregman divergence representation for

the scaled binomial as

$$\begin{aligned}\widetilde{p}(n \mid \widetilde{\eta}, \tilde{\psi}) \equiv \widetilde{p}(n \mid \widetilde{\mu}) &= \exp\{-d_{\widetilde{\phi}}(n, \widetilde{\mu})\} f_{\widetilde{\phi}}(n) \\ &= \exp\{-\beta d_{\phi}(n, \mu)\} f_{\beta\phi}(n).\end{aligned}$$

For any exponential family model we can write its canonical form and construct the corresponding Bregman divergence representations. Thus the same rescaled version can be defined for any exponential family model.

3.2 MAP asymptotics for feature allocations

We use the scaled binomial distribution to develop small-variance asymptotic to the hierarchical model (5). Let $\widetilde{p}_{\beta}(\cdot)$ generically denote distributions under the scaled model. The joint posterior is

$$L(\boldsymbol{Z}, \boldsymbol{w}) \equiv \widetilde{p}_{\beta}(\boldsymbol{Z}, \boldsymbol{w}, \boldsymbol{n} \mid \boldsymbol{N}) = p(\boldsymbol{Z})\, p(\boldsymbol{w} \mid \boldsymbol{Z})\, \widetilde{p}_{\beta}(\boldsymbol{n} \mid \boldsymbol{N}, \boldsymbol{p});$$

based on (6), the scaled binomial likelihood is given by

$$\begin{aligned}\widetilde{p}_{\beta}(\boldsymbol{n} \mid \boldsymbol{N}, \boldsymbol{p}) = \prod_{t=1}^{T} \prod_{s=1}^{S} &\exp\left\{-\beta\left[n_{st} \log(\frac{n_{st}}{N_{st} p_{st}}) + (N_{st} - n_{st}) \log(\frac{N_{st} - n_{st}}{N_{st} - N_{st} p_{st}})\right]\right\} \\ &\times \exp\left\{\beta n_{st}\left[\log(\frac{n_{st}}{N_{st}}) + (N_{st} - n_{st}) \log(\frac{N_{st} - n_{st}}{N_{st}})\right] - h_1(n_{st})\right\},\end{aligned}$$

where $p_{st} = w_{t0} p_0 + \sum_{c=1}^{C} w_{tc} z_{sc}$, as before. Finding the joint MAP of $\boldsymbol{Z}$ and $\boldsymbol{w}$ is equivalent to finding the values of $\boldsymbol{Z}$ and $\boldsymbol{w}$ that minimize $-\log L(\boldsymbol{Z}, \boldsymbol{w})$. We avoid overfitting $\boldsymbol{Z}$ with an inflated number of features by moving the prior towards smaller numbers of features as we increase β. This is achieved by varying $\gamma = \exp(-\beta\lambda^2)$ with increasing β, that is $\gamma \to 0$

as $\beta \to \infty$. Here $\lambda^2 > 0$ is a constant tuning parameter. We show that

$$-\frac{1}{\beta}\log L \;\sim\; \sum_{s=1}^{S}\sum_{t=1}^{T}\{-n_{st}\log(p_{st}) - (N_{st}-n_{st})\log(1-p_{st})\} + C\lambda^2, \tag{9}$$

where C is the random number of columns of $\boldsymbol{Z}$, $u(\beta) \sim v(\beta)$ indicates asymptotic equivalence, i.e., $u(\beta)/v(\beta) \to 1$ as $\beta \to \infty$. The double sum originates from the scaled binomial likelihood and the penalty term arises as the limit of the log IBP prior. The derivation is shown in the Supplementary Material B.

Denote the right hand side of (9) as

$$Q(\boldsymbol{p}) = \sum_{s=1}^{S}\sum_{t=1}^{T}\{-n_{st}\log(p_{st}) - (N_{st}-n_{st})\log(1-p_{st})\} + C\lambda^2, \tag{10}$$

with $p_{st} = w_{t0}p_0 + \sum_{c=1}^{C} w_{tc}z_{sc}$. We refer to (10) as the FL (feature leaning)-means objective function. The first term in the objective functions is a K-means style criterion for the binary matrix $\boldsymbol{Z}$ and $\boldsymbol{w}$ when the number of features is fixed. The second term acts a penalty for the number of selected features. The tuning parameter λ^2 calibrates the penalty. We propose a specific calibration scheme for the application of tumor heterogeneity. See Section 4.2 for details. A local MAP that maximizes the joint posterior $L(\boldsymbol{Z}, \boldsymbol{w})$ is asymptotically equivalent to

$$\text{argmin}_{(C,\boldsymbol{Z},\boldsymbol{w})} Q(\boldsymbol{p}). \tag{11}$$

Keep in mind that $p_{st}(\boldsymbol{Z}, \boldsymbol{w})$ is a function of $\boldsymbol{Z}$ and $\boldsymbol{w}$.

3.3 FL-means algorithm

We develop the FL-means algorithm to solve the optimization problem in (11) and prove its convergence properties.

Since all entries of $\boldsymbol{Z}$ are binary and $w_{tc} \in \mathbb{R}^+$ subject to $\sum_{c=0}^{C} w_{tc} = 1$, (11) is a mixed

integer optimization problem. Mixed integer linear programming (MILP) is NP-hard to solve (Karlof, 2005). The objective function is a non-linear function of $\boldsymbol{Z}$ and $\boldsymbol{w}$, indicating that even sophisticated MILP solvers unlikely benefit this case. Rather than solving the optimization problem (11) as a generic MILP, we construct a coordinate transformation, allowing us to solve this problem as a constrained optimization problem.

Denote $\Delta_{Ct} = \{(w_{t0}, \ldots, w_{tC})^T \in \mathbb{R}^{C+1} \mid \sum_{c=0}^{C} w_{tc} = 1, \text{and } w_{tc} \geq 0 \text{ for all } c\}$. Suppose $\boldsymbol{Z}$ is fixed. Let $\tilde{\boldsymbol{Z}}$ denote $\boldsymbol{Z}$ with the first column multiplied by p_0 and let $H(\boldsymbol{Z}) = \{\tilde{\boldsymbol{Z}}\boldsymbol{w} \mid \boldsymbol{w}_t \in \Delta_{Ct}, t = 1, \ldots, T\}$ denote the set of convex combinations of the column vectors in $\boldsymbol{Z}$ (multiplying the first column by p_0 adds the term $w_{t0}p_0$ in (2)). Then (11) reduces to finding

$$\text{argmin}_{\boldsymbol{p} \in H(\boldsymbol{Z})} Q(\boldsymbol{p}).$$

It can be shown that the objective function $Q(\boldsymbol{p})$ is separable convex (see Supplementary Material C), and this problem can be solved using standard convex optimization methods. However, the next optimization with respect to $\boldsymbol{Z}$ given fixed $\boldsymbol{w}$ is no longer convex. We use a brute-force approach to solve this problem by enumerating all possible $\boldsymbol{Z}$. We propose the following algorithm.

FL-means Algorithm

- Set $C = 1$. Initialize $\boldsymbol{Z}$ as an $S \times (C + 1)$ matrix and assign $z_{s0} = 1$ and $z_{s1} = 1$ with probability 0.5 for $s = 1, \ldots, S$. Initialize $\boldsymbol{w}$ as a $T \times (C+1)$ matrix and $\boldsymbol{w}_t = (w_{t0}, w_{t1})$ are randomly generated from $\text{Dir}(1, 1)$ for $t = 1, \ldots, T$.

- Iterate over the following steps until no changes are made.

    1. For $s = 1, \ldots, S$, optimize (11) with respect to $\boldsymbol{z}_s = (z_{s0}, z_{s1}, \ldots, z_{sC})$, fixing $\boldsymbol{w}_t$ and C at the currently imputed values.

2. For $t = 1, \ldots, T$, optimize (11) with respect to $\boldsymbol{w}_t = (w_{t0}, w_{t1}, \ldots, w_{tC})$ with constraint $\sum_{c=0}^{C} w_{tc} = 1$, fixing $\boldsymbol{Z}$ and C at the currently imputed values.

3. Let $\boldsymbol{Z}'$ equal $\boldsymbol{Z}$ but with one new feature (labeled $C+1$) containing only one randomly selected SNV index s. Set $\boldsymbol{w}'$ that minimizes the objective given $\boldsymbol{Z}'$. If the triplet $(C+1, \boldsymbol{Z}', \boldsymbol{w}')$ lowers the objective (10) from the triplet $(C, \boldsymbol{Z}, \boldsymbol{w})$, replace the latter with the former.

Theorem 3.1. *The FL-means algorithm converges in a finite number of iterations to a local minimum of the FL-means objective function (10).*

See Supplementary Material D for a proof. Theorem 3.1 guarantees convergence, but does not guarantee convergence to the global optimum. From a data analysis perspective the presence of local optima is a feature. It can be exploited to learn about the uncertainty in the posterior distribution by using multiple random initializations. We will demonstrate this use in Sections 4 and 5.

4 Simulation Studies

4.1 Simulation setup and results

We carried out simulation studies to evaluate the proposed FL-means algorithm. We generated a data matrix with $S = 80$ SNVs and $T = 25$ samples. The simulation truth included $C^o = 4$ latent subclones, plus a background subclone ($c = 0$) that included all SNVs. The latent binary matrix $\boldsymbol{Z}^o$ was generated as follows: subclone 1 was characterized by the presence of the first 10 SNVs, subclone 2 by the first 25, subclone 3 by the first 40 and subclone 4 by the first 60. In other words, SNVs 1-10 occurred in all four subclones, SNVs 10-25 in subclones 2-4, SNVs 25-40 in subclones 3-4, SNVs 41-60 in subclone 4 only, and SNVs 61-80 in none of the subclones. Figure 3(a) shows the simulation truth $\boldsymbol{Z}^o$. Let $\pi = (1, 5, 6, 10)$ and

π_p be a random permutation of π, we generated true $\boldsymbol{w}_t^o \sim \text{Dirichlet}(0.2, \pi_p)$, $t = 1, \ldots, T$. Let $p_0^o = 0.01$ and $N_{st} = 50$ for all s and t; we generated $n_{st} \sim \text{Bin}(N_{st}, p_{st})$, where $p_{st} = p_0^o w_{t0}^o + \sum_{c=1}^{C^o} w_{tc}^o z_{sc}^o$. We then ran the FL-means algorithm repeatedly with 1,000 random initializions to obtain a set of local minima of (9) as an approximate representation of posterior uncertainty. Each run of the FL-means algorithm only took 1 minute.

For different λ^2 values, we report point estimates (Figure 3 (b) and (c)) for $\boldsymbol{Z}$ that were obtained by minimizing the objective function in (10). The point estimate is the estimate $(\hat{C}, \hat{\boldsymbol{Z}}, \hat{\boldsymbol{w}})$ that minimizes (10) in 1,000 random initializations. Figure 4 shows the frequencies of the estimated number $\hat{C}$ of features under different λ^2 values (the modes are highlighted by extra large plotting symbols). The plot shows the distribution of local minima over repeated runs of the algorithm with different initializations, all with the same simulated data. The distribution gives some sense of *a posteriori* variability.

An important consideration in the implementation of the FL-means algorithm is the choice of the penalty parameter λ^2. As a sensitivity analysis we ran the FL-means algorithm with different λ^2 values: $\lambda^2 = 2, 4, 6, 8, 10, 20, 200$ and 500. Figure S1 (see Supplementary Material E) shows the estimated number $\widehat{C}$ of features and the realized minimum value of the objective function (10) versus λ^2. We observe that the $\widehat{C}$ decreases and the objective value increases as λ^2 increases. Summarizing Figures 3 and S1, we find that under $\lambda^2 = 8, 10$, both the estimated $\widehat{C}$ and the estimated $\hat{\boldsymbol{Z}}$ perfectly reconstruct the simulation truth $\boldsymbol{Z}^o$. Under $\lambda^2 = 6$, we find $\widehat{C} = 5$ and the estimated $\hat{\boldsymbol{Z}}$ includes true subclones (columns 1-4) as well as an additional spurious subclone that includes some of the SNVs. As expected, a smaller λ^2 leads to a larger $\hat{C}$ value. In summary, the inference summaries are sensitive with respect to the choice of λ^2. A good choice is critical. Below we suggest one reasonable ad-hoc algorithm for the choice of λ^2.

(a) Simulation Truth (b) Estimated (λ^2=10, 8) (c) Estimated (λ^2=6)

Figure 3: A simulation example. Shown are feature allocation matrix $\boldsymbol{Z}$, with shaded area indicating $z_{sc} = 1$, i.e., variant genotype. Rows are SNVs and columns are inferred subclones. Panel (a) displays the simulation truth $\boldsymbol{Z}^o$. Panel (b) displays the estimated feature allocation matrix $\hat{Z}$ when $\lambda^2 = 8$ and 10, which is exactly the same as the simulation truth. Panel (c) displays the estimated feature allocation matrix $\hat{Z}$ when penalty $\lambda^2 = 6$. The first four subclones match the simulation truth.

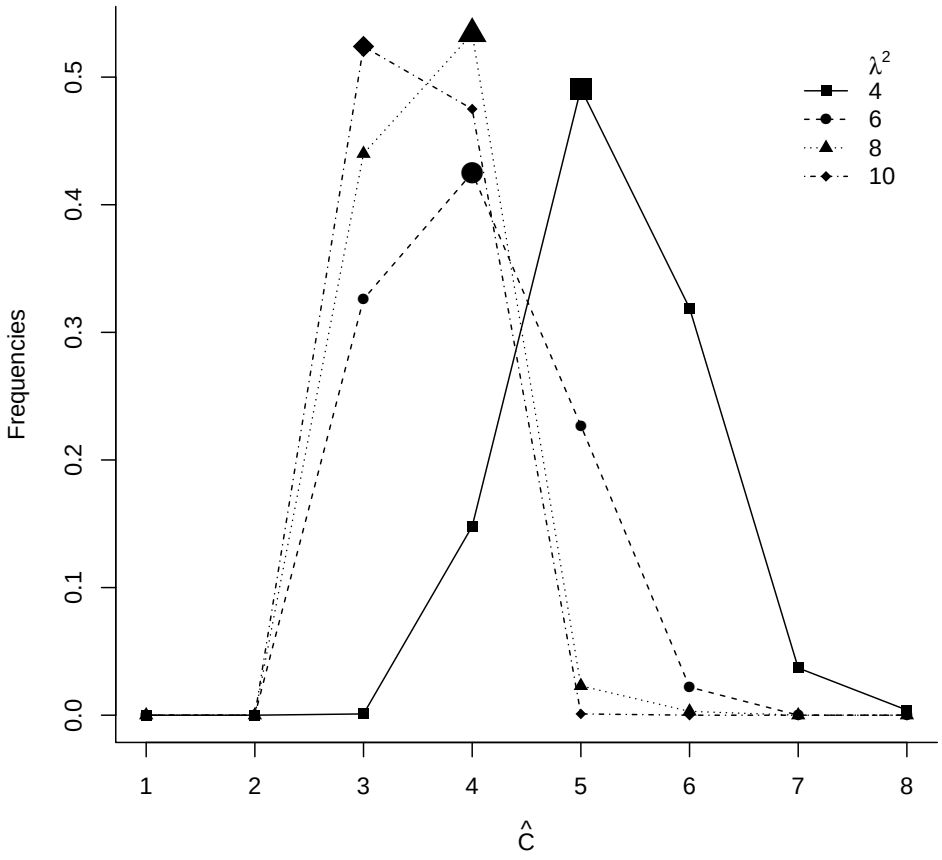

Figure 4: The frequencies of estimated $\hat{C}$ in 1000 simulations with different initializations. The mode is highlighted with extra large plotting symbols.

4.2 Calibration of λ^2

Recall that w_{tc} denotes the relative fraction of subclone c in sample t. Under $\lambda^2 = 8$, the posterior estimate perfectly recovered the simulation truth. The estimates for w_{tc} ranged from 0.01 to 0.53 for $c = 1$, from 0.007 to 0.57 for $c = 2$, from 1.8×10^{-8} to 0.52 for $c = 3$ and from 0.008 to 0.59 for $c = 4$ for $t = 1, \ldots, T$. Posterior inference (correctly) reports that each true subclone c constitutes a substantial part of the composition for some subset of samples. Under $\lambda^2 = 6$, the first four estimated features perfectly recover the simulation truth. However, for $c = 5$, the estimated w_{tc} ranged from 1.1×10^{-9} to 0.06. These very small fractions are biologically meaningless. We find similar patterns under $\lambda^2 = 2$ and 4. Based on these observations, we propose a heuristic to fix the tuning parameter λ^2. We start with a large value of λ^2, say, $\lambda^2 = 50$. While every imputed subclone c constitutes a substantial fraction in some subset of samples, say $w_{tc} > 0.2$ for some t, we decrease λ^2 until newly imputed subclones only take small fraction in all samples, say, $w_{tc} < 1/\hat{C}$.

5 Results

5.1 Pancreatic cancer data

We analyzed exome-sequencing data for five tumor samples from patients with pancreatic ductal adenocarcinoma (PDAC) at NorthShore University HealthSystem (Lee et al., 2013). Since samples were from different patients, we aimed to infer inter-tumor TH. Because individual patients possessed different genomes, we did not assume that the entire genome of a subclone would be shared between samples. Instead, we focused on a set of SNVs, and assumed that the collection of the genotypes at the SNVs could be shared between subclones in different tumor samples. We preselected important SNVs related to PDAC, and performed inference on TH for the local genome. The goal was to reveal if there were subclonal local

genomes shared by the tumor samples.

The mean sequencing depth for the samples was between 60X and 70X. A total of approximately 115,000 somatic SNVs were identified across the five whole exomes using GATK (McKenna et al., 2010). In this subsection we focus on a small number of 118 SNVs selected with the following three criteria: (1) exhibit significant coverage in all samples; (2) occur within genes that are annotated to be related to PDAC in the KEGG pathways database (Kanehisa et al., 2010); (3) are nonsynonymous, i.e., the mutation changes the amino acid sequence that is coded by the gene.

In summary, the PDAC data recorded the total read counts (N_{st}) and variant read counts (n_{st}) of $S = 118$ SNVs from $T = 5$ tumor samples. Figure S2 (see Supplementary Material E) shows the histogram of the observed VAFs, n_{st}/N_{st}. We ran the proposed FL-means algorithm with 1,000 random initializations. Each run of the FL-means algorithm took 50 seconds, while MCMC sampler of finite feature allocation model proposed by Lee et al. (2013) took over one hour for each fixed number of subclones. After searching for reasonable values of λ^2 using the algorithm from section 4.2, we estimated $\hat{C}$ to be 5, i.e., five subclones.

Figure 5 summarizes the analysis results. In (a) we present an estimate of $\boldsymbol{Z}$, as the most frequently converged value in our runs. Each of the five columns represents a subclone and each row represent an SNV, with green and red blocks indicating wildtype and mutant genotypes. In (b) we use a circos plot (Krzywinski et al., 2009) to present the proportions of the five estimated subclones for each sample. Samples 1 and 5 possess the same three subclones (2, 4, and 5) while samples 2-4 all possess subclone 1 and another distinct sublcone. The results show that each tumor sample consists of multiple subclones, exhibiting somatic mosaicism. Tumor samples are heterogeneous; some samples have subclones with the same local genomes (SNVs), such as samples 1 and 5, and other samples possess their own unique local genomes, such as samples 2-4. Even though samples 1 and 5 possess the same subclones, their proportions vary. Lastly, subclones 1 , 4, and 5 are more prevalent than the other

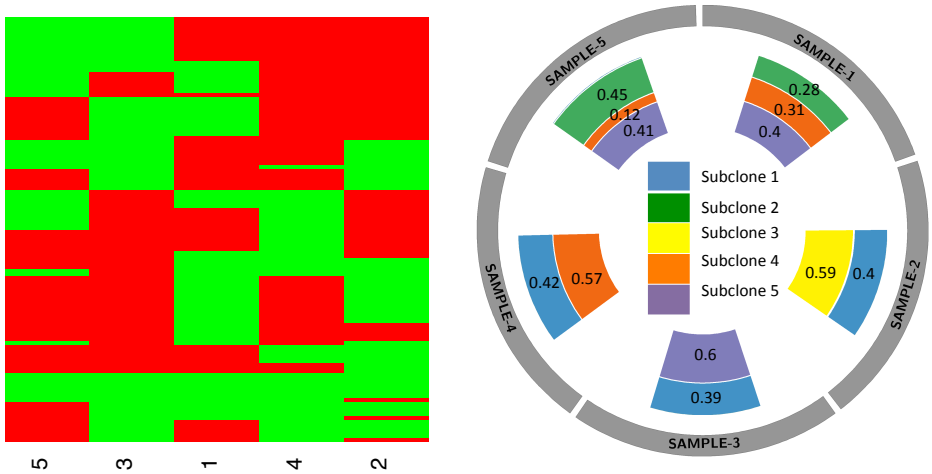

(a) (b)

Figure 5: Summary of TH analysis using the PDAC data with 118 SNVs. (a): Posterior estimated subclonal genotypes (columns) of the selected SNVs. Five subclones are identified. A green/red block indicates a mutant/wildtype allele at the SNV in the subclone. (b): A circos plot showing the estimated proportion $\hat{w}_{tc}$ with $\hat{C} = 5$ subclones.

subclones, both appearing in three out of five samples. Subclone 3 is the least prevalent appearing only in one sample.

Examining closely the Z matrix in Figure 5, we do not see any intra-sample clonal evolution pattern. That is, we do not see a phylogenic relationship between any subclones in the same sample. This is not surprising since the subclones within a sample could be decedents of different parent subclones in that sample which are not observed.

The shared local genomes, or the SNVs can be further annotated for functional investigation. We did not pursue further as such an investigation is beyond the scope of this paper. Nevertheless, the information presented in Figure 5 will set an example for the type of inference one can achieve using the proposed methodology.

Lastly, to examine the gain in the computation using the proposed method, we re-analyzed the PDAC data, but now keeping all SNVs that exhibit significant coverage in all samples and occur in at least two samples– not limited to those in KEGG pathways.

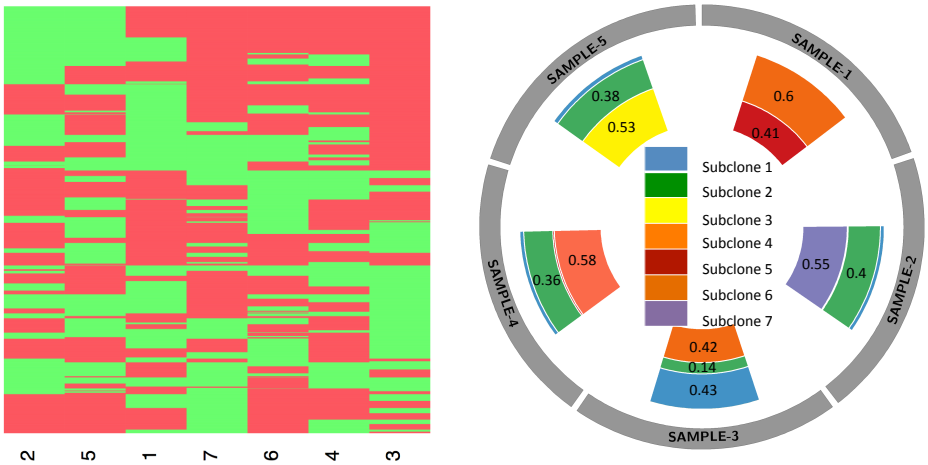

(a) (b)

Figure 6: Summary of TH analysis using the PDAC data with 6,599 SNVs. (a): Posterior estimated subclonal genotypes (columns) of the selected SNVs. Five subclones are identified. A green/red block indicates a mutant/wildtype allele at the SNV in the subclone. (b): A circos plot showing the estimated proportion $\hat{w}_{tc}$ with $\hat{C} = 7$ subclones.

This filtering left us with $S = 6,599$ SNVs. With such a large data size, in practice it is infeasible to run the MCMC sampler proposed in Lee et al. (2013). We applied the proposed FL-means algorithm with 1,000 random initializations. Each run of the FL-means algorithm took less than 2 minutes. After searching for λ^2 using the suggested heuristic, we estimated $\hat{C}$ to be 7. The full inference results are summarized in Figure 6. Interestingly, we found similar patterns to the previous analysis, such as that each sample possesses multiple subclones and samples 2-4 share a common subclone while possessing additional distinct ones. However, with more SNVs, there are now three (as opposed to one) distinct subclones, each only appearing in one out of five samples. This is not surprising since with more somatic SNVs, by definition the chance of having common subclones between different tumor samples will reduce.

5.2 Breast cancer data

As the second example, we analyzed Illumina HiSeq data from Horvath et al. (2013) who obtained variant signatures of 17 breast cancer patient tissues using whole transcriptome RNA-sequencing. We downloaded the data from Gene Expression Omnibus (GEO) (Barrett et al., 2013). The 17 samples consist of six triple negative breast cancer samples (TNBC), which lack expression of three therapeutically significant components: estrogen receptor (ER), progesterone receptor (PR) and the Human Epidermal Growth Factor Receptor 2 (HER2); six Non-TNBC (ER, PR, and HER2-positive); and five HER2-positive samples (ER and PR negative).

For each sample, we kept SNVs with reads coverage ≥ 100 and VAFs between 0.2 and 0.8. The total number of candidate SNVs after combining the filtered SNVs from all samples was 6,992. Finally, we restricted our attention to SNVs occurring in at least three tumor samples, which left us with $S = 329$ SNVs.

We identified $\hat{C} = 4$ different subclones for these 17 breast cancer samples. Figure 7 shows $\hat{\boldsymbol{w}}$, arranged by the three clinical subtypes (knowledge of the subtypes was not used in the model). Samples of different clinical subtypes show distinct patterns of subclone proportions, although subclones 1 and 3 are prevalent in most samples. More specifically, the HER2 samples consist of relatively constant proportions of subclones 2 and 4, while most non-TNBC samples show a depletion of subclone 2 and most TNBC samples show a depletion of both subclones 2 and 4. Each of the subclones is characterized by unique SNV profiles. For example, subclone 1 harbors more SNVs (194) than other subclones and many of these SNVs (69) are unique to subclone 1. Noteworthy, some of these unique SNVs carried by subclone 1 are associated with breast cancer, such as chr13:38138689:C/T:1, which changes an amino acid *valine* to *methionine* in the protein *periostin*. Periostin is a component of the extracellular matrix that promotes metastasis formation by increasing Wnt signaling (Malanchi et al., 2012) and is predicted to be damaging by both, the Sift algorithm (Kumar

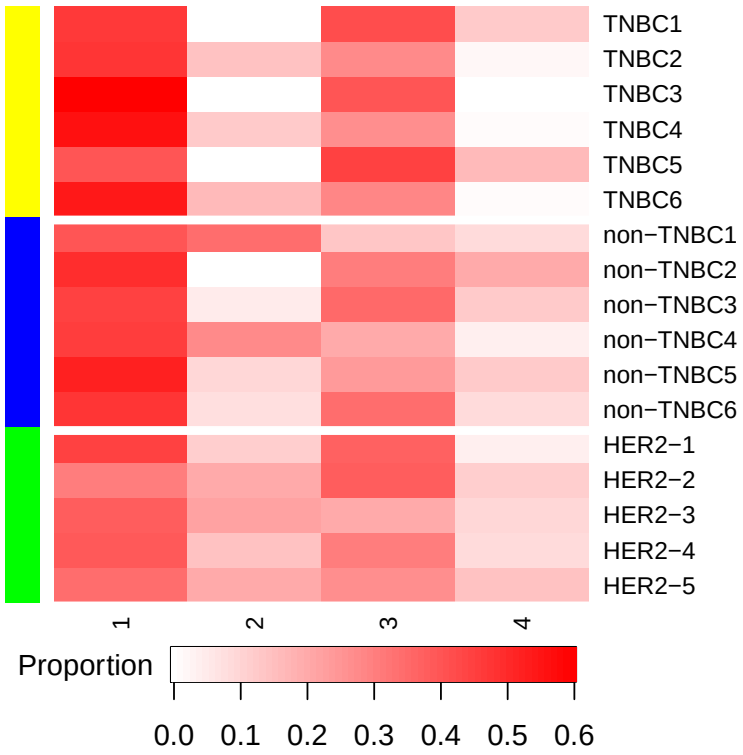

Figure 7: Summary of TH analysis using the breast cancer data. Shown is a heatmap of the estimated subclonal proportions $\hat{w}_{tc}$ with $\hat{C} = 4$ subclones.

et al., 2009) and the Polyphen algorithm (Adzhubei et al., 2010).

6 Conclusion

We introduce a small-variance asymptotic analysis for the MAP in a feature allocation model under a binomial likelihood as it arises in inference for tumor heterogeneity. The proposed FL-algorithm uses a scaled version of the binomial likelihood which can be introduced as a special case of scaled exponential family models that are based on writing the sampling models in terms of Bregman divergences. The algorithm provides simple and scalable inference to feature allocation problems.

The discussion moves beyond the Gaussian distribution setting and opens the door to new applications of feature allocation models, including applications involving discrete data. Although the proposed algorithm is developed and illustrated for binomial sampling model, it can be easily extended to any exponential family model. The approach included a pragmatic method for selecting the tuning parameter λ^2, based on the implied inference on the number of features C.

Inference in the three datasets that we analyzed, including the two PDAC data sets and the breast cancer data, indicates that subclones with shared local genomes do recur across different patients. Such sharing would provide valuable information for future disease prognosis, taking advantage of the innovation proposed in this paper. For example, Ding et al. (2012) demonstrated that the relapse of acute myeloid leukemia was associated with new mutations acquired by a subpopulation of cancer cells derived from the original population.

Supplementary Material

Supplementary material is available on http://www.ma.utexas.edu/users/yxu/.

Acknowledgment

Peter Müller and Yuan Ji's research is partially supported by NIH R01 CA132897.

References

Adzhubei, I. A., Schmidt, S., Peshkin, L., Ramensky, V. E., Gerasimova, A., Bork, P., Kondrashov, A. S., and Sunyaev, S. R. (2010). A method and server for predicting damaging missense mutations. *Nature methods*, 7(4):248–249.

Banerjee, A., Merugu, S., Dhillon, I. S., and Ghosh, J. (2005). Clustering with bregman divergences. *The Journal of Machine Learning Research*, 6:1705–1749.

Barrett, T., Wilhite, S. E., Ledoux, P., Evangelista, C., Kim, I. F., Tomashevsky, M., Marshall, K. A., Phillippy, K. H., Sherman, P. M., Holko, M., et al. (2013). Ncbi geo: archive for functional genomics data sets?update. *Nucleic acids research*, 41(D1):D991–D995.

Bregman, L. M. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217.

Broderick, T., Jordan, M. I., and Pitman, J. (2012a). Beta processes, stick-breaking and power laws. *Bayesian analysis*, 7(2):439–476.

Broderick, T., Kulis, B., and Jordan, M. I. (2012b). Mad-bayes: Map-based asymptotic derivations from bayes. *arXiv preprint arXiv:1212.2126*.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38.

Ding, L., Ley, T. J., Larson, D. E., Miller, C. A., Koboldt, D. C., Welch, J. S., Ritchey, J. K., Young, M. A., Lamprecht, T., McLellan, M. D., et al. (2012). Clonal evolution in relapsed acute myeloid leukaemia revealed by whole-genome sequencing. *Nature*, 481(7382):506–510.

Gerlinger, M., Rowan, A. J., Horswell, S., Larkin, J., Endesfelder, D., Gronroos, E., Martinez, P., Matthews, N., Stewart, A., Tarpey, P., et al. (2012). Intratumor heterogeneity and branched evolution revealed by multiregion sequencing. *New England Journal of Medicine*, 366(10):883–892.

Ghoshal, S. (2010). The Dirichlet process, related priors and posterior asymptotics. In Nils Lid Hjort, Chris Holmes, P. M. and Walker, S. G., editors, *Bayesian Nonparametrics*, pages 22–34. Cambridge University Press.

Green, P. J. (1995). Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika*, 82(4):711–732.

Griffiths, T. L. and Ghahramani, Z. (2006). Infinite latent feature models and the Indian buffet process. In Weiss, Y., Schölkopf, B., and Platt, J., editors, *Advances in Neural Information Processing Systems 18*, pages 475–482. MIT Press, Cambridge, MA.

Hartigan, J. A. and Wong, M. A. (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108.

Hastie, T., Tibshirani, R., and Friedman, J. J. H. (2001). *The elements of statistical learning*, volume 1. Springer New York.

Horvath, A., Pakala, S. B., Mudvari, P., Reddy, S. D. N., Ohshiro, K., Casimiro, S., Pires, R., Fuqua, S. A., Toi, M., Costa, L., et al. (2013). Novel insights into breast cancer genetic variance through rna sequencing. *Scientific reports*, 3.

Kanehisa, M., Goto, S., Furumichi, M., Tanabe, M., and Hirakawa, M. (2010). Kegg for representation and analysis of molecular networks involving diseases and drugs. *Nucleic acids research*, 38(suppl 1):D355–D360.

Karlof, J. K. (2005). *Integer programming: theory and practice*. CRC Press.

Krzywinski, M., Schein, J., Birol, I., Connors, J., Gascoyne, R., Horsman, D., Jones, S., and Marra, M. (2009). Circos: an information aesthetic for comparative genomics. *Genome Res.*, 19(9):1639–45.

Kulis, B. and Jordan, M. I. (2011). Revisiting k-means: New algorithms via bayesian nonparametrics. *arXiv preprint arXiv:1111.0352*.

Kumar, P., Henikoff, S., and Ng, P. C. (2009). Predicting the effects of coding non-synonymous variants on protein function using the sift algorithm. *Nature protocols*, 4(7):1073–1081.

Landau, D. A., Carter, S. L., Stojanov, P., McKenna, A., Stevenson, K., Lawrence, M. S., Sougnez, C., Stewart, C., Sivachenko, A., Wang, L., et al. (2013). Evolution and impact of subclonal mutations in chronic lymphocytic leukemia. *Cell*, 152(4):714–726.

Lee, J., Müller, P., Ji, Y., and Gulukota, K. (2013). A feature allocation model for tumor heterogeneity. *Technical Report*.

Liu, J. S. (2008). *Monte Carlo strategies in scientific computing*. springer.

Malanchi, I., Santamaria-Martínez, A., Susanto, E., Peng, H., Lehr, H.-A., Delaloye, J.-F., and Huelsken, J. (2012). Interactions between cancer stem cells and their niche govern metastatic colonization. *Nature*, 481(7379):85–89.

McKenna, A., Hanna, M., Banks, E., Sivachenko, A., Cibulskis, K., Kernytsky, A., Garimella, K., Altshuler, D., Gabriel, S., Daly, M., and Mark, D. (2010). The genome analysis toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. *Genome research*, 20(9):1297–1303.

Teh, Y. W., Görür, D., and Ghahramani, Z. (2007). Stick-breaking construction for the indian buffet process. In *International Conference on Artificial Intelligence and Statistics*, pages 556–563.

Thibaux, R. and Jordan, M. (2007). Hierarchical beta processes and the indian buffet process.

In *Proceedings of the 11th Conference on Artificial Intelligence and Statistics (AISTAT)*, Puerto Rico.